arXiv:1603.00652v2 [cs.RO] 2 Dec 2016

Physics-Based Damage-Aware Manipulation Strategy Planning Using Scene Dynamics Anticipation

Tobias Fromm and Andreas Birk

***Abstract*— We present a damage-aware planning approach which determines the best sequence to manipulate a number of objects in a scene. This works on task-planning level, abstracts from motion planning and anticipates the dynamics of the scene using a physics simulation. Instead of avoiding interaction with the environment, we take unintended motion of other objects into account and plan manipulation sequences which minimize the potential damage. Our method can also be used as a validation measure to judge planned motions for their feasibility in terms of damage avoidance. We evaluate our approach on one industrial scenario (autonomous container unloading) and one retail scenario (shelf replenishment).**

I. Introduction

Since no robot is perfect, just like the humans building them, autonomous manipulation can be destructive if fragile goods have to be handled. Many compliant grippers, sophisticated perception systems and manipulation routines have been developed recently, mostly relying on consequent obstacle avoidance. This works to a certain degree, but why not take alteration of the environment into account for planning instead of bluntly avoiding interaction with it?

We present a manipulation strategy planning approach which validates and selects the best sequence to remove or unload a number of objects from a scene while taking into account possible damage-prone movements of any other objects. The idea relies on anticipating the dynamics of the scene using a physics simulation which contains a copy of the physical scene, including the robot, the environment and a number of objects to manipulate.

The possibilities are obvious: One the one hand, motion planning is facilitated because of a less constrained search space in contrast to regarding all other objects as obstacles. In some scenes, it may be hard or even impossible to perform manipulation if no collisions are permitted. On the other hand, we want to avoid damage to the possibly heavy or fragile goods as well as the robot itself which may occur if an object is shifted or dropped unintendedly.

Hence, the goal of our method is to optimize a robot's autonomous behavior according to the anticipated dynamics of the real-life scene. This works without imposing explicit spatial or logical dependencies between objects, thus the approach is not tied to a particular application domain.

The authors are with the Robotics Group, Computer Science & Electrical Engineering, Jacobs University Bremen, Germany; *t.fromm@jacobs-university.de*. The research leading to the presented results has received funding from the European Union's Seventh Framework and Horizon 2020 programs within the projects "Cognitive Robot for Automation Logistics Processes" (RobLog) and "Effective Dexterous ROV Operations in Presence of Communication Latencies" (DexROV).

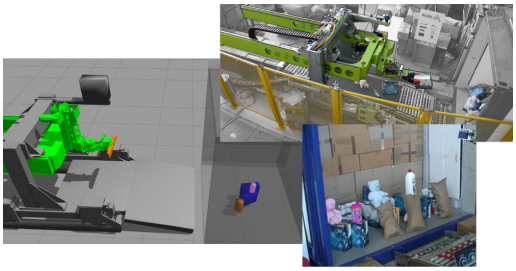

(a) Logistics (container unloading) [1]

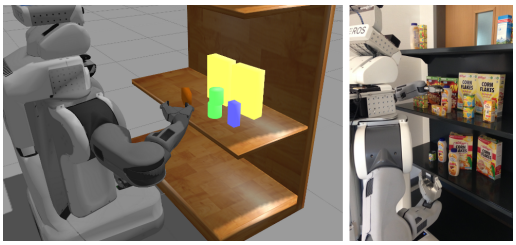

(b) Supermarket (shelf replenishment) [2]

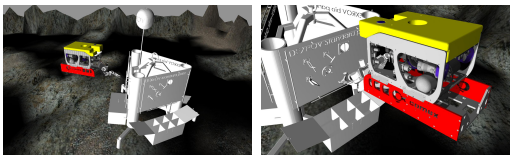

(c) Underwater (maintenance, archaeology and biology) [3]

Fig. 1. Typical application scenarios for autonomous manipulation

We evaluate our approach on two everyday scenarios shown in Figure 1: logistics, in the context of the EU project "Cognitive Robot for Automation of Logistics Processes" (RobLog) [1] where containers are unloaded autonomously, and in a supermarket for shelf replenishment [2].

Since our method is generally applicable for robotic manipulation tasks, underwater environments like used within the EU project "Effective Dexterous ROV Operations in Presence of Communication Latencies" (DexROV) [3] may serve as another application example.

The contribution of our work is:

- a method to plan strategies for the manipulation of a sequence of objects
- minimizing unintended object movement while dynamically considering the physics of all objects
- including an industrial and a retail application scenario which show that our approach can be utilized in arbitrary usage domains.

II. Foundations and Related Work

A. Physics-Based Motion Planning

Motion planning based on physics simulation has been investigated in several publications, amongst these, Zickler and Veloso's work [4]. Instead of conventionally planning motions for a robot's physical degrees of freedom, they group low-level motions and use them as a single point in action selection sampling space. As for the environment representation, Zickler and Veloso state that it is necessary to define domain-appropriate distance functions between states of the scene in order to be able to compare them. For their usage domain of robot soccer and minigolf, the Euclidean distance serves as a sufficient measure in most cases.

Weitnauer et al. [5] present an approach which works in a similar way, though only on plain objects and using only the Euclidean distance as a distance function.

Another usage example of physics-based planning can be found in [6] and [7] where the authors present methods of planning grasps through clutter, taking shifting objects into account and explicitly manipulating them. This is rather similar to our approach, though not with the intention to find an efficient manipulation order, but to clear manipulation paths from obstructing objects.

Before using one of these approaches, the object to manipulate would have to be selected by some algorithm or a human operator. On the contrary, our method automates exactly this scenario where not only physics-based manipulation planning takes place, but also autonomous object selection to plan manipulation sequences on a high level.

In addition, in the scenarios intended to be tackled with our approach, intentionally pushing objects aside with the gripper poses an additional risk of damaging the goods and even the gripper. For this reason, we want to limit the movement of any other than the object to manipulate as much as possible. Our approach takes care of this by penalizing any movement of a *passive object*, i.e. an object which is not subject to intended manipulation at the moment. Therefore, intentional as well as unintentional damage-prone manipulation of passive objects is avoided.

B. Manipulation Strategy Planning

In contrast to the methods described so far, our approach does not aim on planning *motions*, but high-level *behaviors* with the help of a physics simulation. This means that, instead of giving an explicit trajectory, our method will a) *validate* whether an ongoing manipulation is likely to be successful or b) *plan* a sequence of manipulation actions to perform in the desired scenario. Therefore, our possible application scenarios are separated into *Manipulation Outcome Prediction* and *Manipulation Sequence Planning* which will be introduced in the following.

a) Manipulation Outcome Prediction: Covered by this term, several methods have shown different ideas and scenarios. Unfortunately, most of the existing work focuses on predicting the outcome of motions instead of behaviors, like Pastor et al.'s work [8] where robot parameters like joint positions and velocities are continuously monitored and verified whether they lie within a learned envelope. Haidu et al. [9] use a similar technique, also to learn motoric skills.

In contrast to these works, Rockel et al. [10] present an interesting idea of predicting task success on the basis of motion of the manipulated objects. Concretely, their application includes a robot which maximizes its speed moving around in a scene, but without losing the balance of an object which sits on its base. This works on the basis of semantic predicates, but with hand-crafted features which describe if the object shakes or topples and which cannot easily be employed on a different scenario.

Akhtar et al. [11] also use semantic predicates to reason about success of a manipulation action. They present a Machine Learning algorithm which is able to predict whether an object released on top of another will behave as expected. This detection of external faults can predict the behavior of a concrete action based on a simulated training set, but again uses hand-crafted predicates which only cover this particular action. Our approach, on the contrary, aims on covering implicit manipulation of passive objects as well, independent of which exact manipulation is performed on the scene.

b) Manipulation Sequence Planning: In order to manipulate objects which are placed currently unreachable behind others, Stilman et al. [12] use a sampling-based planner to move away the blocking objects.

Okada et al. [13] present a similar simple, but effective example for strategy planning for a humanoid robot where obstacles are considered as planning goals and subsequently moved away from the desired walking path. Their goal is to enable robots to implicitly perform necessary manipulation actions, even if these actions were not part of the original task plan. This is similar to our recent work evaluated in a different scenario [2], but different to what we present in this work. Here we use a dynamics simulation to evaluate the feasibility of a manipulation action, taking unintended side effects into account in addition to static spatial knowledge.

III. Prerequisites

Simulating the behavior of robots and objects requires both high accuracy and some abstractions since not every detail in robot and object properties can easily be modeled. Since our approach is intended to work in complex scenarios based on perception systems for object recognition and real robots for motion execution, we need to be sure that these work error-free and take care of the respective noise where it occurs. Therefore, we assume object models and environment to be given as well as the state of the robot.

A. Simulation Scene Composition

In order to build the scene configuration used in the physics simulation, we first recognize and localize all objects in the scene using our perception system [1] [14]. This system includes a pipeline of several segmentation and filtering steps before it executes diverse recognition modules, including a feature-based textured object recognition module [15] and a graph-based shape model recognition module [16].

B. Motion Planning, Grasp Planning and Execution

After all objects have been perceived, the respective grasps need to be planned, evaluated and finally executed by the robot. The way of determining grasping configurations strongly depends on the objects and gripper used in the scenario since different object sizes and structure demand for different gripper sizes and amounts of dexterity. Since the focus in this work does not lie on grasp and motion planning and our method is supposed to be re-used in different scenarios, we use simple grasping configurations around the principal object axes.

The physics simulation uses a kinematics and dynamics model of the robot which has been developed together with the respective simulated controllers. For the supermarket scenario, we use the standard PR2 simulation model. The eventual generation of the motion trajectories can be performed, for example, by using *MoveIt!*.

IV. Manipulation Cost Functions

A. Motivation

The usage of cost functions to find the most suitable and reasonable solution to a planning problem has been common practice for a long time (e.g. in [13], [17]) to be able to evaluate the effects of a particular action. In order to apply this method to a specific planning problem, domain knowledge has to come into play which accounts for the optimization target of the respective problem.

Since our approach aims on reducing unintended motions of *passive objects*, i.e. objects which are not subject to intended manipulation at the moment, we need to consider cost functions which cover spatial modifications in the scene.

B. Terminology

For all of the following definitions, we use $\mathcal{O}$ as the set of objects present in the scene, $\alpha \in \mathcal{O}$ as the *active object* (which will be manipulated), $\Phi = \mathcal{O} \setminus \alpha$ as the set of *passive objects* (which will not be manipulated) and $\phi \in \Phi$ as one member of this set.

C. Pose-Based Cost Functions

The following two simple cost functions are based on object poses and allow for a quick validation without a lot of processing, but may not be predictive-efficient enough.

a) Maximum Pose Shift c_p: describes the maximum Euclidean distance between object poses:

$$c_p = \max_{\phi \in \Phi} \|\mathbf{p}_{t+1}(\phi) - \mathbf{p}_t(\phi)\| \tag{1}$$

where $\mathbf{p}_t(\phi) \in \mathbb{R}^3$ is the position of passive object ϕ before (at time t) and $\mathbf{p}_{t+1}(\phi)$ after running the simulation step (at time $t+1$). This is simple and easy to compute, but smoothes out non-straight motion paths and thus does not describe changes in an object's direction of movement.

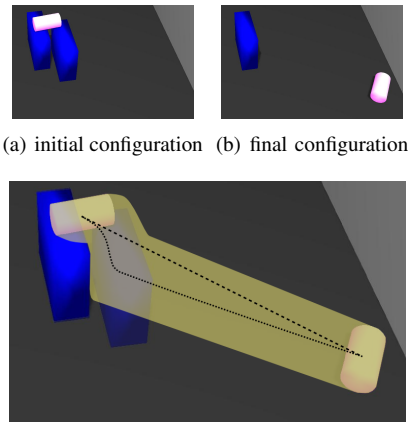

(a) initial configuration (b) final configuration

(c) Pose Shift (dashed line), Path Length (dotted line) and Swept Convex Volume (yellow)

Fig. 2. Visualization of the Swept Convex Volume compared to the Pose Shift and Path Length covered by a passive object during manipulation. When the right box was pulled out by the robot, the cylindrical container fell down and rolled away, causing costs like shown.

b) Maximum Path Length c_l: is the maximum length of the path any object's centroid traveled during simulation:

$$c_l = \max_{\phi \in \Phi} \sum_{i=1}^{n} \|\mathbf{p}_i(\phi) - \mathbf{p}_{i-1}(\phi)\| \tag{2}$$

where $\mathbf{p}_i(\phi) \in \mathbb{R}^3$ is a position of passive object ϕ during the simulation step (= between time steps t and $t+1$) and n is the number of positions covered during simulation. This takes into account different directions of movement, but not a possible spin the object may have been exposed to.

D. Volume-Based Cost Functions

The following cost functions are based on the volume of the objects instead of their pose, thus considering complex motion paths as well as changes of movement direction and spin.

a) Maximum Swept Convex Volume c_v: Swept Volume estimations can be used for collision detection [18] as well as space occupancy calculation for automation and production purposes [19]. For a mathematical formulation see the survey of Abdel-Malek et al. [20].

The object's surface mesh is used as the *generator* which follows a *trajectory* (a set of poses covered while moving during simulation runtime) and creates the *swept volume* (the outer boundary of the object during its motion). This can be considered a voxel-based variant of [20] which is widely used in applications where generator and trajectory are discrete.

In most applications, this volume basically represents a concave hull around all points in space which the object ever touched. However, this takes a lot of effort to compute [19] and, since the concave hull of a number of points is not generally well-defined, may raise ambiguities for different kinds of objects. As a remedy, we simplified the original Swept Volume approach by replacing the concave with the convex hull, normalized by object volume, which is easy to compute and well-defined, and call the result the *Swept*

Convex Volume. Figure 2 shows a visualization of the Swept Convex Volume of a moving object compared to its Pose Shift and Path Length.

Since we relate all objects in a scene, we take the *Maximum Swept Convex Volume* c_v which is the maximum over all Swept Convex Volumes of the scene objects. It is determined as

$$c_v = \max_{\phi \in \Phi} V(\phi) \tag{3}$$

where $V(\phi)$ is the Swept Convex Volume of the respective object $\phi \in \Phi$ as computed in Algorithm 1.

Algorithm 1 Swept Convex Volume calculation

1: **input**: object mesh $\mathcal{M}(\phi)$, object poses $\mathbf{p}_{0..n}(\phi)$ covered during simulation
2: create point cloud $\mathcal{C}(\phi)$ from $\mathcal{M}(\phi)$ at $\mathbf{p}_0(\phi)$
3: **for all** $\mathbf{p}_i(\phi)$ **do**
4: $\quad \mathcal{C}_i(\phi) \leftarrow \mathcal{C}(\phi)$ transformed from $\mathbf{p}_0(\phi)$ to $\mathbf{p}_i(\phi)$
5: $\quad \mathcal{C}(\phi) \leftarrow \mathcal{C}(\phi) \cup \mathcal{C}_i(\phi)$
6: **end for**
7: $H(\phi) \leftarrow \text{convhull}(\mathcal{C}(\phi))$
8: $H_0(\phi) \leftarrow \text{convhull}(\mathcal{C}_0(\phi))$
9: $V(\phi) \leftarrow \text{volume}(H(\phi)) \ / \ \text{volume}(H_0(\phi))$
10: **output**: Swept Convex Volume $V(\phi)$

b) Maximum Weighted Swept Convex Volume c_w: Since the desired usage of our method puts special focus on fragile and vulnerable goods, we propose an additional cost function based on c_v, but with additional weights for each of the 6-D axes. These weights can be adapted to the domain, e.g. to punish vertically dropping objects like when they fall off a shelf or out of a container. Additionally, one could think of applications like objects running on a conveyor where a lateral translation could push the goods off the belt, thus stopping them to be conveyed. Rotational weights may be employed in a domain dealing with objects which spill when they are turned, for example liquid containers or open underwater objects like amphoras and treasure chests.

As for the implementation of such weights, in Algorithm 1 we need to replace $\mathbf{p}_i(\phi)$ in Line 4 with the weighted $\mathbf{p}_i^w(\phi)$:

$$\mathbf{p}_i^w(\phi) = \mathbf{p}_0(\phi) + \text{diag}(\mathbf{w}) \cdot (\mathbf{p}_i(\phi) - \mathbf{p}_0(\phi)) \tag{4}$$

where $\mathbf{w} = \begin{bmatrix} w_x & w_y & w_z & w_\varphi & w_\theta & w_\psi \end{bmatrix}^\intercal$ are domain-dependent weights for each of the components of the 6-D object pose.

V. Manipulation Outcome Prediction

Scene Dynamics Anticipation as a validation method for planned actions can be used in runtime-restricted scenarios and scenarios where another planner has already determined which objects to manipulate in which order. One example for this is the planner presented by Mojtahedzadeh et al. [21] which uses static equilibrium calculations to find objects that support each other, hence it will eventually select objects first which do not support any other object.

Implicitly, our approach will usually determine similar manipulation sequences, but in addition considers dynamic events occuring during the manipulation for which Mojtahedzadeh et al.'s approach is not equipped. Consequently, using our method for validating plans on dynamic scenarios can be seen as an addition and enrichment to another planner.

Algorithm 2 shows how we perform Outcome Prediction using Scene Dynamics Anticipation with the definitions and cost functions given in the previous sections.

Algorithm 2 Manipulation Outcome Prediction

1: **input**: scene objects $\mathcal{O}$, active object $\alpha \in \mathcal{O}$
2: spawn all objects $\mathcal{O}$ in simulation
3: plan approach T_1^α and extract T_2^α trajectory
4: move simulated robot on T_1^α
5: grasp active object α
6: move simulated robot on T_2^α
7: release active object α
8: $c \leftarrow$ manipulation costs ($\rightarrow$ Section IV)
9: **output**: manipulation costs c

One drawback of using Scene Dynamics Anticipation as a validation method, however, is the fact that it outputs the manipulation costs for the given scene configuration, but classification between positively and negatively validated actions has to be performed manually via a threshold.

This is a common problem for Outcome Prediction methods in general, also for Rockel et al.'s approach [10] where thresholds are defined manually for which an object is considered to be toppling. Another example is [8] where the authors perform statistical tests on whether a planned motion lies within a learned envelope. Whenever such an approach is deployed on a new scenario, these parameters have to be set accordingly.

We empirically determined the manipulation cost threshold based on satisfactory operation of our method. For the logistics scenario, for instance, where the container is low above ground and provides lots of operation space, a threshold of $c_w = 2.0$ gives satisfactory results. The implicit lower limit for this value is $c_w = 1.0$ which equals the volume of the object. Since jitter can always occur in simulation and small pushes usually do not cause any damage, setting this value too close to 1.0 will prevent the algorithm from positively evaluating any configuration. On the other hand, too loose settings should be avoided.

VI. Manipulation Sequence Planning

In addition to Manipulation Outcome Prediction, we now want to introduce a way to include Scene Dynamics Anticipation into planning a sequence of manipulation actions.

Instead of serving as a validation tool only, our method is able to plan a manipulation sequence which considers side-effects of moving objects. In addition, obstacle-avoiding motion planning may fail in certain scenarios. But since we explicitly allow for moving passive objects during manipulation, though trying to minimize it, our high-level planner may still succeed in finding a viable manipulation sequence.

In our previous work [1], we have sketched the difficulties and pitfalls of motion planning in heavily confined spaces,

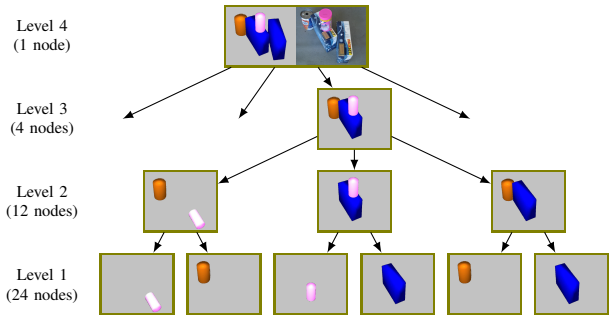

Fig. 3. Example search tree (other branches cropped for better visibility) – each node shows its initial configuration, i.e. before manipulation

Algorithm 3 Search Tree Generation for Sequence Planning

initialize search tree $\mathcal{S}_0 \leftarrow \emptyset$
initialize set of objects $\mathcal{O}_0$ with current scene
procedure CREATENODE($\mathcal{S}_i$, $\mathcal{O}_i$)
  if $|\mathcal{O}_i| \leq 1$, i.e. this is a leaf node **then**
    return
  else
    create new tree node N_i containing object set $\mathcal{O}_i$
    $\mathcal{S}_{i+1} \leftarrow \mathcal{S}_i \cup N_i$
    determine new active object $\alpha_i \in \mathcal{O}_i$
    $\mathcal{O}_{i+1} \leftarrow \mathcal{O}_i \setminus \alpha_i$
    return CREATENODE($\mathcal{S}_{i+1}$, $\mathcal{O}_{i+1}$)
  end if
end procedure
output: filled search tree $\mathcal{S}$

i.e. when all passive objects are considered as obstacles. This may even lead to a dead end in manipulating any remaining object if the constraints are too tight. Manipulation Sequence Planning using Scene Dynamics Anticipation, on the other hand, allows for passive objects to move and takes their motion into account as a feature, thus getting stuck in such a case is less likely.

A. Planning Procedure

Our method uses a *search tree* built from all possible objects configurations which can occur while manipulating all objects in the scene, see the example in Figure 3. The initial configuration is shown as the tree root, on the top, while further down the tree one object has been unloaded in each node. The images show the respective initial configuration for these nodes, before performing the manipulation step.

While traversing the tree further down, unintended motion and high costs are caused for the pink container which falls off the blue box when this one is manipulated. The leaves merely contain one remaining object which can directly be manipulated disregarding the scene physics since no passive objects are left.

Algorithm 3 shows how to generate the search tree for a scene given the initial object configuration. After the tree has been filled, traversing it works like shown in Algorithm 4, using a depth-first search-like technique and exploiting our Manipulation Outcome Prediction method (Algorithm 2).

Algorithm 4 Manipulation Sequence Planning

input: scene objects $\mathcal{O}$, search tree $\mathcal{S}$ generated from $\mathcal{O}$ ($\rightarrow$ Algorithm 3)
for all nodes $N_i \in \mathcal{S}$ **do**
  perform Outcome Prediction ($\rightarrow$ Algorithm 2)
end for
for all leaf nodes $N_j^{\text{leaf}} \in \mathcal{S}$ **do**
  $c_j \leftarrow$ summed-up costs of N_j^{leaf}'s parents
  $c_{\min} \leftarrow \min(c_{\min}, c_j)$
end for
output: node N_j with lowest manipulation costs $c_{\min}$

As for the cost functions used for calculating the node costs, all those introduced in Section IV can in principle be used interchangeably here. Since one of the main topics of this paper is damage avoidance, however, we prefer the Maximum Weighted Swept Convex Volume c_w because it takes into account a) complex motion paths and b) change of movement direction and spin. Both of these factors are a crucial ingredient to avoid objects toppling or dropping under all circumstances.

In order to use A^* or a similar search algorithm for our problem we would need to design an admissible heuristic. However, with the cost function used in our method it is not possible to define such a heuristic because the distance (in costs) to the target configuration cannot be determined.

B. Efficiency Considerations

In contrast to A^* or other common search methods like discretized Rapidly-Exploring Random Trees (RRT) or Rapidly-Exploring Random Leafy Trees (RRLT) [22], our scenario causes more of an efficiency issue in estimating the node costs than in the search itself. Traversing the tree taken by itself is rather trivial and fast because the number of nodes which are handled in our search trees can be processed rapidly on modern machines. Thus, the efficient search problem is overshaded by a heuristic cost estimation problem during the processing of each individual node which is prominent and crucial to solve in minimal time.

The *worst-case* tree size for a number of objects n, constructed from all possible permutations of the object set like in Algorithm 3, is

$$\sum_{i=1}^{n} \frac{n!}{i!} = n! + \frac{n!}{2} + \ldots + \frac{n!}{(n-1)!} + 1 \qquad (5)$$

where $\frac{n!}{i!}$ is the number of objects for the respective tree level i, $i = 1$ representing the bottom level (i.e. leaves) and $i = n$ the top level (i.e. root); see the example in Figure 3.

This number grows fast with a rising number of objects, taking along the cost estimation problem with the same speed. A possible remedy for this is to partition the scene into piles of 4-5 objects and plan a manipulation sequence for each of these piles individually. However, how to optimize these partitions still remains an open research question.

In addition to not using more than 4-5 objects at a time, we cut off as many branches of the search tree as possible in

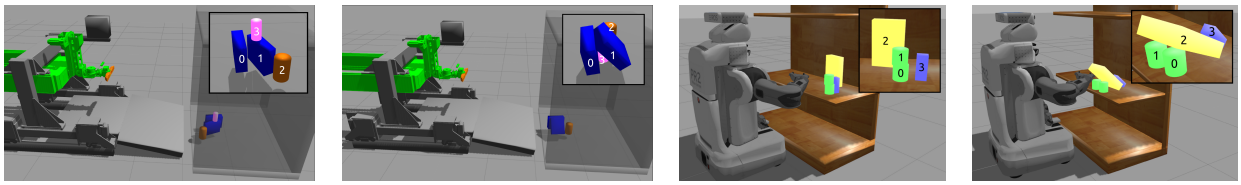

(a) Scene 1 (b) Scene 2 (c) Scene 3 (d) Scene 4

Fig. 4. Experimental scenes used for evaluating our method

which ultimately there is no possibility to present the optimal solution. In the following, we will present several ideas to reduce the tree size as far as possible. In Section VII we show that the tree size in a typical scenario can be reduced to as low as 48.7% using these measures.

1) Implicit Search Tree Pruning: As we will show by an empiric consideration in the Results section, in every scene there is a number of configurations which implicitly cannot be simulated. Usual reasons for this are that approaching an object may push a passive object which, in turn, moves the active object away by a significant distance. In this case, the active object will end up unreachable for the planned manipulation action. Additionally, sometimes no feasible grasping configuration can be found without the robot colliding with the environment (container, shelf, etc.). Since motion planning is out of scope of our work, we have to skip simulating the respective configuration in this case and impose infinite costs.

2) Explicit Search Tree Pruning: In addition to the implicit reasons given above, the used depth-first search allows for cropping tree branches: After simulating a configuration, if the accumulated costs in the current branch exceed the total costs of any already computed goal sequence (i.e. a tree leaf), we stop exploring the current branch.

3) Re-Use of Similar Configurations: Another way to reduce computational load, in addition to pruning tree branches, is to re-use similar configurations which occured before somewhere else in the tree. Before running the simulation, we compare the current configuration to each one which was simulated before, anywhere in the tree, for similarity of objects and their poses. If there is a similar configuration, we re-use this node and the whole child tree of the node (if any) without having to simulate any of them.

Generally speaking, apart from the mentioned measures to reduce computational complexity, the runtimes of our method strongly depend on the used simulator. Moreover, regarding the total runtime of one sequence planning run, motion planning and execution take the biggest part of time which is very robot and application-specific.

Fortunately, our method is well-suited for parallelization because all nodes on the same tree level can be simulated in parallel since their configurations do not depend on each other. As the number of objects in the scene and, with it, the number of tree levels rises, the limit for the speedup achieved by parallelization is the number of used threads / CPU cores.

C. Open-Loop vs. Closed-Loop Planning

The method we propose basically allows for open-loop planning, where the manipulation sequence is planned once and then executed without further anticipation of scene dynamics, or closed-loop, where the whole process is repeated from the beginning after manipulation of any object. Replanning the entire remaining manipulation sequence after each perception/manipulation cycle in a closed loop has the advantage of covering object movement which stems from to inaccurate physics simulation or other, external disturbances in the scene. However, time constraints may not allow for running Manipulation Sequence Planning after each manipulation cycle, so using the results only as an initial plan is a viable alternative.

VII. Applications and Experimental Results

In this section, we want to evaluate our method applied on two of the described practical scenarios. Typical scenes for both are shown in Figure 1.

As for the first scenario, logistics is a field where many different goods have to be handled, some of which are heavy, bulky, fragile or otherwise damage-prone. Our previous work on unloading shipping containers [1] provides an ideal application in this respect, so we have modeled the robot, container and objects from this scenario as our first example.

Secondly, domestic and retail robotics provide another playground for manipulation strategy planning, thus the supermarket scenario of [2] will serve as the second example.

A. Prerequisites

For evaluating our method, we take several modalities as granted and thus abstract from them because they fall out of the scope of this work.

First of all, since the perception and grasping accuracy of our previous work have been validated extensively already in [1], [14], [15] and [16], we are not going into detail on these measures anymore.

Abstracting from details in the object models is necessary as well because we would like to show the general concept of our method in this work. Although it is desirable to obtain simulation models as close as possible to the real objects, this constitutes an own branch of research. Because of this, we use a simplified representation of the rigid objects from our real scenarios (see Figure 1) like shown in Figure 4. For the same reason, we do not employ complex grasping techniques here, but instead generate simple grasping configurations around the principal axes.

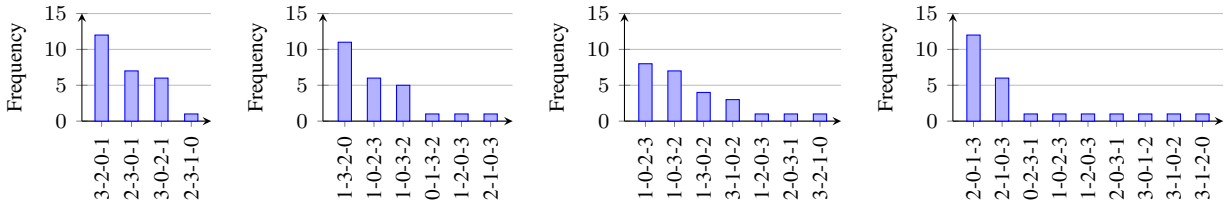

(a) Scene 1 (b) Scene 2 (c) Scene 3 (d) Scene 4

Fig. 5. Frequencies of first-ranked manipulation sequences selected by our method

Although we present a general approach which is not limited to scenarios where damage-awareness is the main focus, we want to emphasize the importance of this factor in our applications. Thus, we use the Maximum Weighted Swept Convex Volume c_w as a cost function which explicitly penalizes covered space. The weights $\mathbf{w}$ for c_w we set on

$$\mathbf{w} = \begin{bmatrix} 1.0 & 1.0 & 2.0 & 1.0 & 1.0 & 1.0 \end{bmatrix}^\mathsf{T}$$

which puts increased importance on damage-prone vertical motion. This is desired especially in scenarios like a supermarket where objects would break if dropped from a shelf.

B. Experimental Scenes

We want to show how our method works by running it many times on different scenes and comparing the final manipulation sequences. Additionally, we show what happens to the search tree which can be heavily pruned during traversal.

a) Logistics: Scenes 1 and 2 in Figure 4 show typical scenarios encountered in RobLog [1] where Scene 2 is the more challenging one due to the goods supporting each other. Hence, once any of the objects is manipulated it is likely that one of the others will move as well.

b) Supermarket: The tall shelf, along with the cans which may roll away if dropped, is a hostile area for any damage-prone product and thus a good example for our method. The tipped-over items in Scene 4 occur frequently in a supermarket where customers leave the shelf like this, products coming to rest partly on top of each other.

C. Manipulation Sequence Planning Results

We ran our method 25 times each on the two scenes per application, resulting in a total of 100 runs. Figure 5 shows a histogram of how often a particular manipulation sequence was selected for a particular scene of Figure 4.

Numerical results are shown in Table I with their means (μ) and standard deviations (σ) for the following criteria:

- *mean first-ranked costs per node*: mean over 25 runs of the mean costs per object imposed on the manipulation sequence selected by Algorithm 4
- *mean second-ranked costs per node*: mean over 25 runs of the mean costs per object imposed on the second-best manipulation sequence
- *pruned tree nodes*: percentage of nodes pruned from the tree, resulting in a similar reduction in runtime, and composed of the following sub-criteria:
  - *known subtree*: a node shares its object configuration and poses with a previously processed node, thus the costs were copied without re-simulating
  - *costs exceed existing sequence*: a solution is existing already which has lower costs than the costs accumulated so far in this branch
  - *active object moved*: the active object was pushed away during approach, ending up unreachable
  - *object out of workspace*: an object fell out of the workspace (container/shelf), causing maximum damage and ending up unreachable for the robot
  - *planning failure*: no motion plan was found for the active object, e.g. because it was pushed away too far in a parent node
- *nodes with significant movement*: percentage of nodes which were not pruned from the tree and reported significant costs higher than a manually defined threshold

D. Discussion

The results in Figure 5 show that, over many runs, of the possible number of sequences only a small selection is considered as best. The fact that not the same sequence is reported for each run, but there is obviously some noise in the selection stems mostly from the used motion planning techniques. These rely on random trees and thus deliver noisy results for similar inputs. This way, different grasping configurations are selected for the same object configuration, resulting in different motions even if the scene looks similar.

Nevertheless, the resulting manipulation sequences are valid and reasonable in a way that they comply with the intuitive sequence a human would use to clear the container/shelf: unloading the objects first which do not physically support others. So, even if there is some distribution of the finally selected manipulation sequence due to reasons that cannot be influenced by our planning approach, our method succeeds in selecting reasonable solutions for all scenes.

The results show that generally a major ratio of pruned nodes results from comparing the current costs to the already planned sequences. Therefore, this is one of the most useful tricks we use to reduce computational load.

In Scene 3, items falling off the shelf are the major reason for tree pruning: 31.9% of nodes featured an object which moved out of the robot's workspace. The robot used in the logistics scenario has a big workspace in all directions, additionally, the container provides physical limits on three sides, thus this did not happen for Scenes 1 and 2.

TABLE I

NUMERICAL RESULTS

Scene	**1**		**2**		**3**		**4**		**Total**	
	μ	σ	μ	σ	μ	σ	μ	σ	μ	σ
mean first-ranked costs per node	1.008	0.008	1.208	0.074	1.349	0.359	1.975	0.321	1.382	0.435
mean second-ranked costs per node	1.337	0.868	1.429	0.326	1.700	1.728	3.764	7.393	2.051	3.878
pruned tree nodes	34.8%		50.4%		58.3%		61.9%		51.3%	
- known subtree	7.5%	4.8%	2.6%	2.5%	1.2%	2.2%	0.0%	0.0%	2.8%	4.1%
- costs exceed existing sequence	16.5%	8.3%	18.3%	5.3%	19.3%	14.7%	41.0%	18.1%	23.8%	16.0%
- active object moved	5.8%	5.6%	25.2%	13.5%	4.2%	11.6%	8.1%	9.2%	10.9%	13.2%
- object out of workspace	0.0%	0.0%	0.0%	0.0%	31.9%	15.7%	7.0%	9.6%	9.6%	15.9%
- planning failure	5.0%	9.6%	4.3%	8.4%	1.7%	3.1%	5.8%	10.9%	4.2%	8.6%
for reference: non-pruned nodes with significant movement ($c_w > 2.0$)	13.8%	3.0%	11.9%	4.9%	14.8%	4.3%	18.9%	5.1%	14.9%	5.0%

The high values for the standard deviations of pruned nodes result from all the children of such a node automatically being cut off. This way, any disturbance propagates down through the tree and may affect many nodes at the same time when it happens in a high-level node.

VIII. CONCLUSION

We presented a method to plan manipulation sequences for a number of objects, depending on their physical behavior. This enables autonomous robots to anticipate an optimal manipulation order to avoid potential damage by taking into account unintended movement of other objects.

Our approach presents merely a first attempt to solve a complex problem. However, since the planning technique is abstract, given the scene was perceived correctly, it is universally usable on any autonomous robot scenario.

We evaluated our approach on two everyday scenarios showing good results for scenes of different physical complexity and on different types of robots.

REFERENCES

[1] T. Stoyanov, N. Vaskevicius, C. A. Mueller, T. Fromm, R. Krug, V. Tincani, R. Mojtahedzadeh, S. Kunaschk, R. Mortensen Ernits, D. Ricao Canelhas, M. Bonilla, S. Schwertfeger, M. Bonini, H. Halfar, K. Pathak, M. Rohde, G. Fantoni, A. Bicchi, A. Birk, A. Lilienthal, and W. Echelmeyer, "No More Heavy Lifting: Robotic Solutions to the Container Unloading Problem," *Robotics and Automation Magazine*, 2016, in press.

[2] J. Winkler, F. Balint-Benczedi, T. Wiedemeyer, M. Beetz, N. Vaskevicius, C. A. Mueller, T. Fromm, and A. Birk, "Knowledge-Enabled Robotic Agents for Shelf Replenishment in Cluttered Retail Environments," in *International Conference on Autonomous Agents and Multiagent Systems*, 2016.

[3] J. Gancet, P. Weiss, G. Antonelli, M. Pfingsthorn, S. Calinon, A. Turetta, C. Walen, D. Urbina, S. Govindaraj, P. Letier, X. Martinez, J. Salini, B. Chemisky, G. Indiveri, G. Casalino, P. di Lillo, E. Simetti, D. de Palma, A. Birk, T. Fromm, C. Mueller, A. Tanwani, I. Havoutis, A. Caffaz, and L. Guilpain, "Dexterous Undersea Interventions with Far Distance Onshore Supervision: the DexROV Project," in *IFAC Conference on Control Applications in Marine Systems*, 2016.

[4] S. Zickler and M. Veloso, "Efficient Physics-Based Planning: Sampling Search Via Non-Deterministic Tactics and Skills," in *International Conference on Autonomous Agents and Multiagent Systems*, 2009.

[5] E. Weitnauer, R. Haschke, and H. Ritter, "Evaluating a Physics Engine as an Ingredient for Physical Reasoning," in *International Conference on Simulation, Modeling, and Programming for Autonomous Robots*, 2010.

[6] M. Dogar, K. Hsiao, M. Ciocarlie, and S. Srinivasa, "Physics-Based Grasp Planning Through Clutter," in *Robotics: Science and Systems*, 2012.

[7] N. Kitaev, I. Mordatch, S. Patil, and P. Abbeel, "Physics-Based Trajectory Optimization for Grasping in Cluttered Environments," in *International Conference of Robotics and Automation*, 2015.

[8] P. Pastor, M. Kalakrishnan, S. Chitta, E. Theodorou, and S. Schaal, "Skill learning and task outcome prediction for manipulation," in *International Conference on Robotics and Automation*, 2011.

[9] A. Haidu, D. Kohlsdorf, and M. Beetz, "Learning Action Failure Models from Interactive Physics-based Simulations," in *International Conference on Intelligent Robots and Systems*, 2015.

[10] S. Rockel, S. Konecny, S. Stock, J. Hertzberg, F. Pecora, and J. Zhang, "Integrating Physics-Based Prediction with Semantic Plan Execution Monitoring," in *International Conference on Intelligent Robots and Systems*, 2015.

[11] N. Akhtar, A. Kuestenmacher, P. Ploeger, and G. Lakemeyer, "Simulation-based approach for avoiding external faults," in *International Conference on Advanced Robotics*, 2013.

[12] M. Stilman, J.-U. Schamburek, J. Kuffner, and T. Asfour, "Manipulation Planning Among Movable Obstacles," in *International Conference on Robotics and Automation*, 2007.

[13] K. Okada, A. Haneda, H. Nakai, M. Inaba, and H. Inoue, "Environment Manipulation Planner for Humanoid Robots Using Task Graph That Generates Action Sequence," in *International Conference on Intelligent Robots and Systems*, 2004.

[14] N. Vaskevicius, C. Mueller, M. Bonilla, V. Tincani, T. Stoyanov, G. Fantoni, K. Pathak, A. Lilienthal, A. Bicchi, and A. Birk, "Object Recognition and Localization for Robust Grasping with a Dexterous Gripper in the Context of Container Unloading," in *International Conference on Automation Science and Engineering*, 2014.

[15] N. Vaskevicius, K. Pathak, A. Ichim, and A. Birk, "The Jacobs Robotics approach to object recognition and localization in the context of the ICRA'11 Solutions in Perception Challenge," in *International Conference on Robotics and Automation*, 2012.

[16] C. Mueller, K. Pathak, and A. Birk, "Object Shape Categorization in RGBD Images using Hierarchical Graph Constellation Models based on Unsupervisedly Learned Shape Parts described by a Set of Shape Specificity Levels," in *International Conference on Intelligent Robots and Systems*, 2014.

[17] K. Goldberg, "Stochastic Plans for Robotic Manipulation," Ph.D. dissertation, 1990.

[18] H. Taeubig, B. Baeuml, and U. Frese, "Real-time Swept Volume and Distance Computation for Self Collision Detection," in *International Conference on Intelligent Robots and Systems*, 2011.

[19] A. von Dziegielewski, M. Hemmer, and E. Schoemer, "High Precision Conservative Surface Mesh Generation for Swept Volumes," *Transactions on Automation Science and Engineering*, vol. 12, no. 1, pp. 764–769, 2015.

[20] K. Abdel-Malek, J. Yang, D. Blackmore, and K. Joy, "Swept Volumes: Fundation [*sic*], Perspectives, and Applications," *International Journal of Shape Modeling*, vol. 12, no. 1, 2006.

[21] R. Mojtahedzadeh, A. Bouguerra, E. Schaffernicht, and A. Lilienthal, "Support relation analysis and decision making for safe robotic manipulation tasks," *Robotics and Autonomous Systems*, vol. 71, no. July 2015, pp. 99–117, 2015.

[22] S. Morgan and M. Branicky, "Sampling-Based Planning for Discrete Spaces," in *International Conference on Intelligent Robots and Systems*, 2004.